

Fiscal Year 2022

Ver. 2023-01-19a

Course number: CSC.T433  
School of Computing,  
Graduate major in Computer Science

# Advanced Computer Architecture

## 9. Instruction Level Parallelism: Out-of-order Execution and Multithreading



[www.arch.cs.titech.ac.jp/lecture/ACA/](http://www.arch.cs.titech.ac.jp/lecture/ACA/)  
Room No.W831, **HyFlex**  
Mon 13:45-15:25, Thr 13:45-15:25

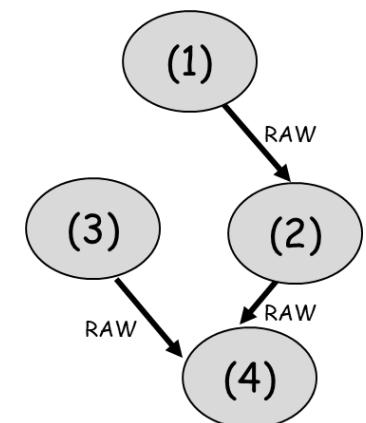
Kenji Kise, Department of Computer Science  
kise\_at\_c.titech.ac.jp

# Exploiting Instruction Level parallelism (ILP)

- A superscalar has to handle some flows efficiently to exploit ILP
  - Control flow (control dependence)
    - To execute  $n$  instructions per clock cycle, the processor has to fetch at least  $n$  instructions per cycle.
    - The main obstacles are branch instruction (BNE, BEQ)
    - Prediction
    - Another obstacle is instruction cache
  - Register data flow (data dependence)
    - Out-of-order execution
      - Register renaming
      - Dynamic scheduling
    - Memory data flow
      - Out-of-order execution
      - Another obstacle is data cache

(1) add \$5,\$1,\$2  
(2) add \$9,\$5,\$3  
(3) lw \$4, 4(\$7)  
(4) add \$8,\$9,\$4

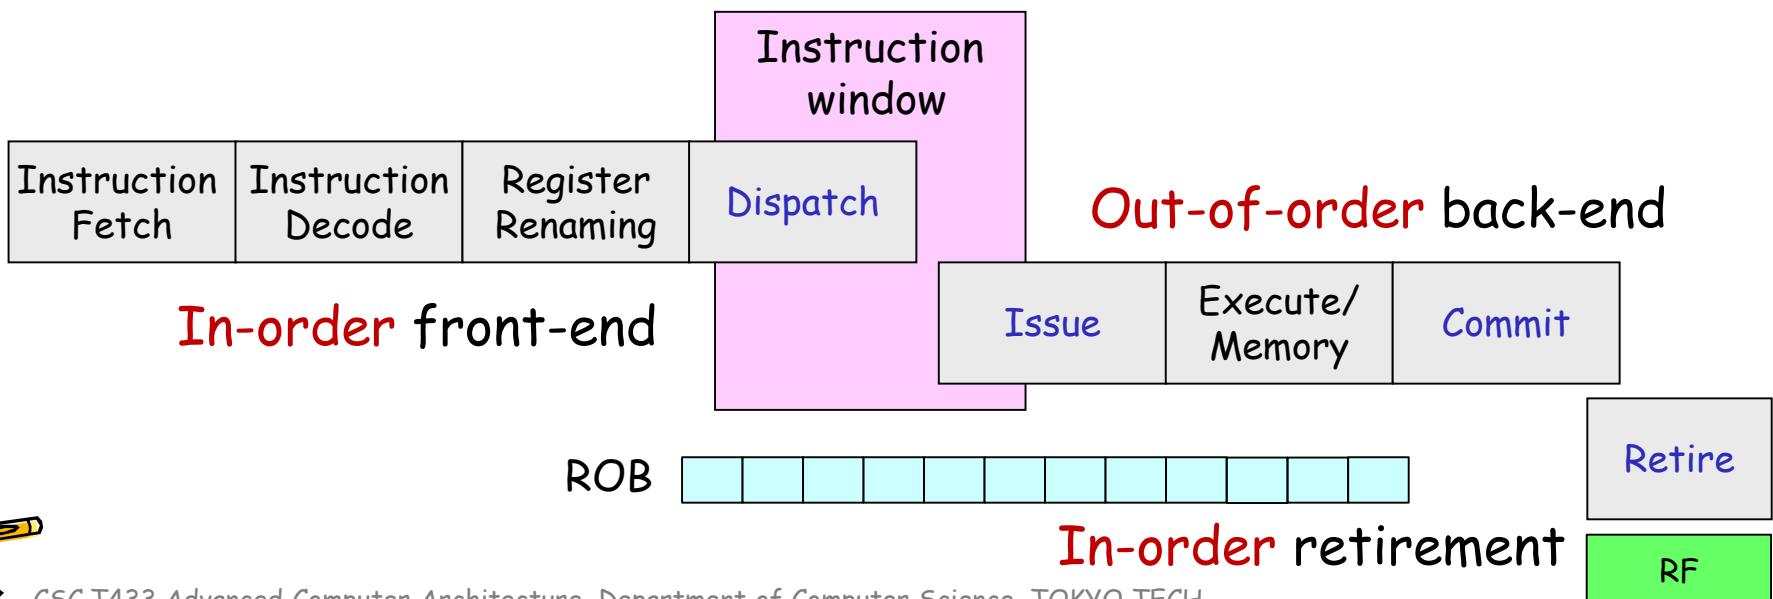
(3) lw \$4, 4(\$7)  
(1) add \$5,\$1,\$2  
(2) add \$9,\$5,\$3  
(4) add \$8,\$9,\$4



# Instruction pipeline of OoO execution processor

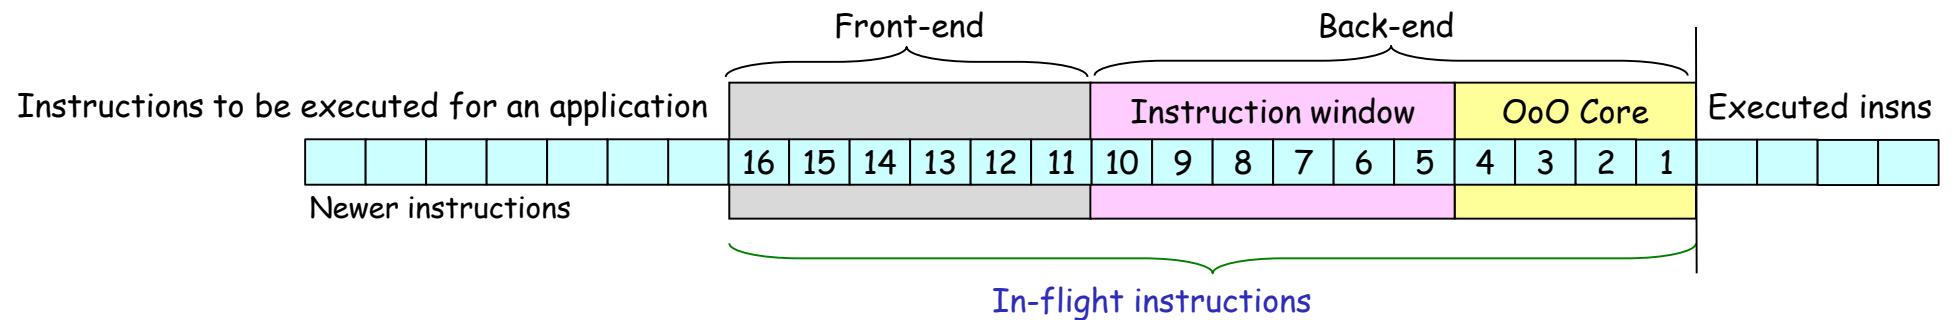
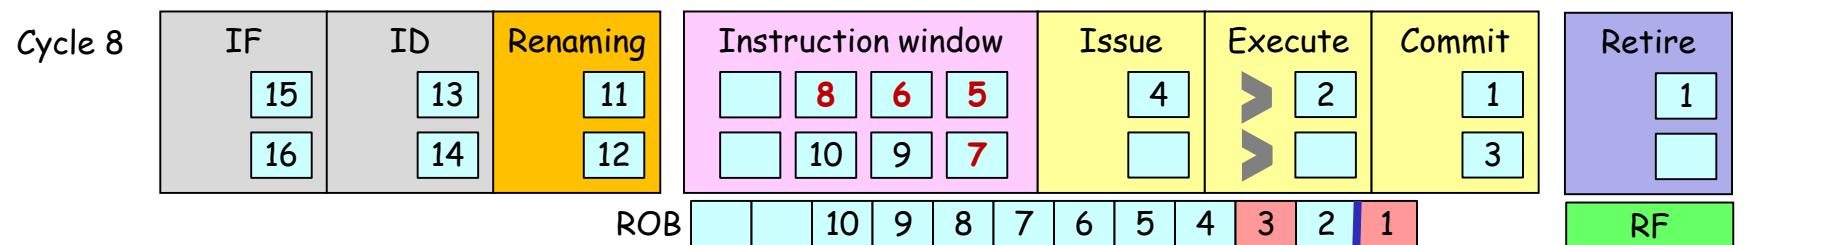
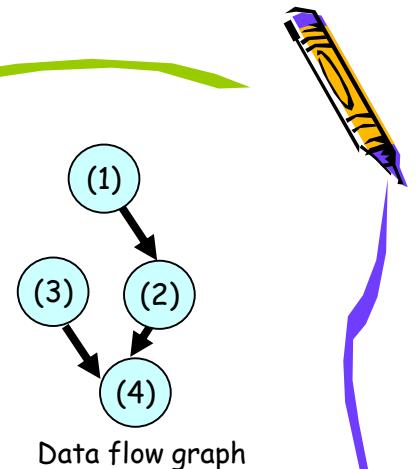


- Allocating instructions to **instruction window** is called **dispatch**
- Issue** or **fire** wakes up instructions and their executions begin
- In **commit** stage, the computed values are written back to **ROB** (**reorder buffer**)
- The last stage is called **retire** or **graduate**. The completed **consecutive** instructions can be retired.  
The result is written back to **register file (architectural register file of 32 registers)** using a logical register number from \$0 to \$31.



# Register dataflow

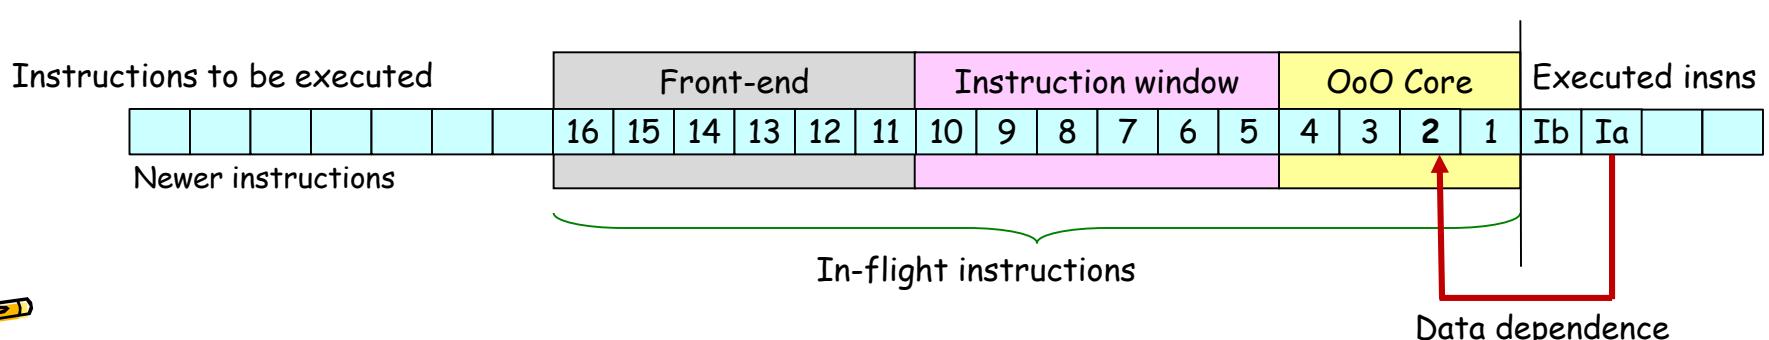
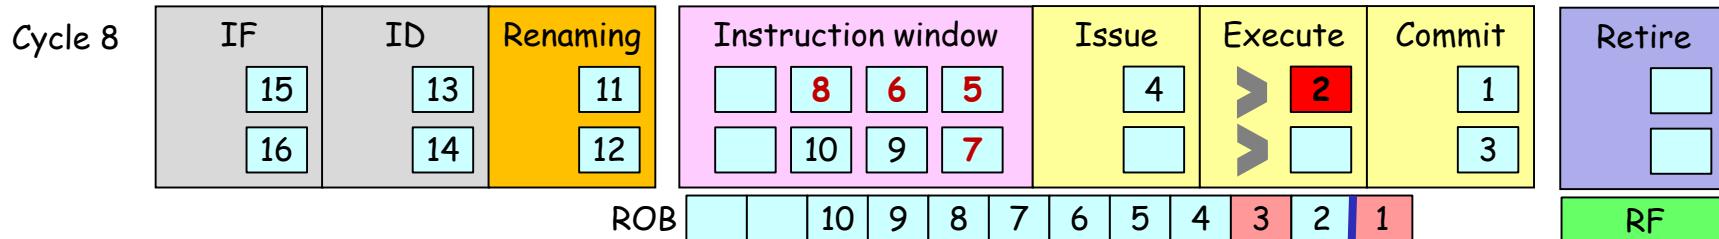
- **In-flight instructions** are ones processing in a processor



# Case 1: Register dataflow from a far previous insn

- One source operand of insn I2 is from a retired instruction Ia.
- Because Ia is retired long ago, the physical destination register has been freed. The tag of the source register \$3 can not be renamed at the renaming stage for I2, still having the logical register tag \$3.
- Where does the operand \$3 of I2 come from?

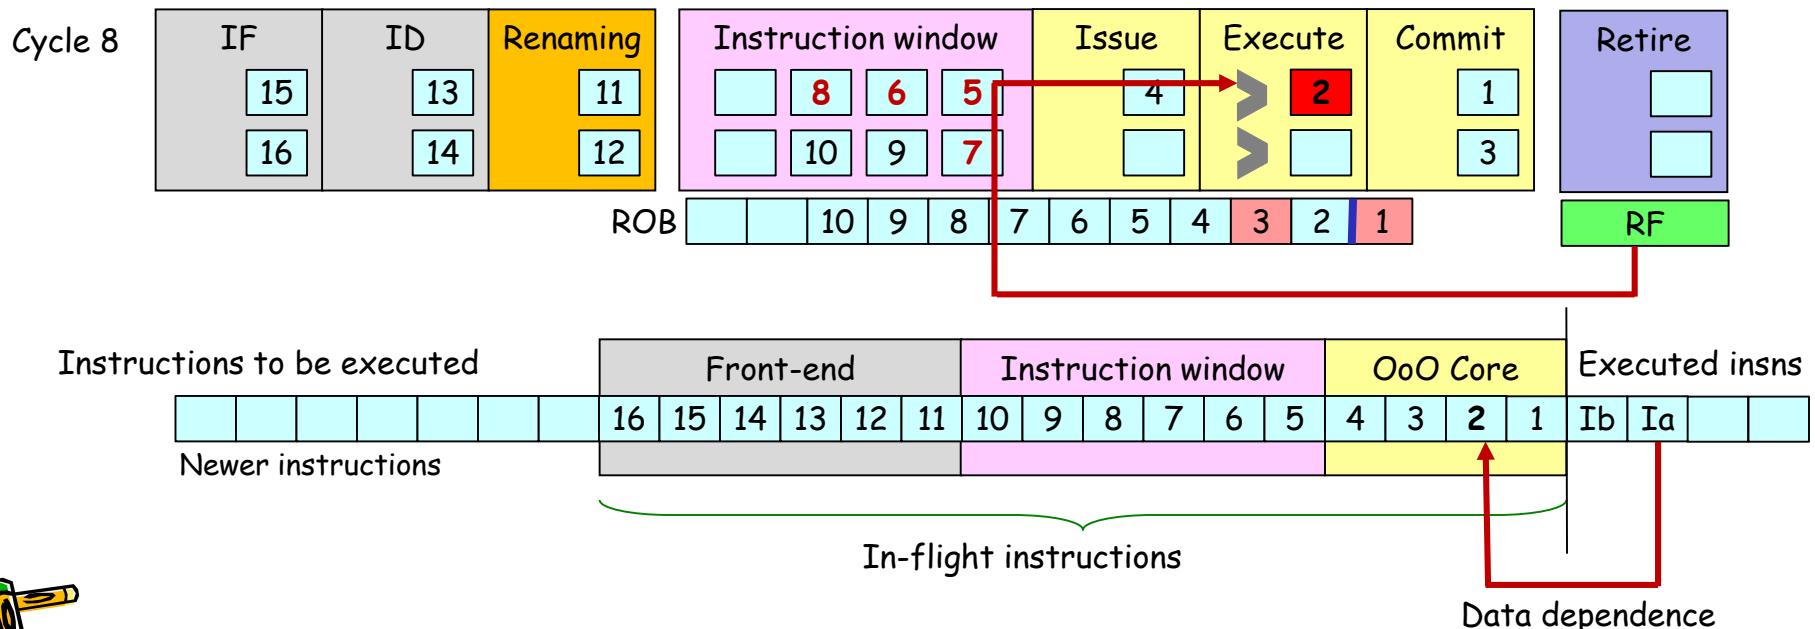
Ia: add \$3,\$0,\$0  
 I1: sub p9,\$1,\$2  
 I2: add p10,p9,\$3  
 I3: or p11,\$4,\$5  
 I4: and p12,p10,p11



# Case 1: Register dataflow from a far previous insn

- One source operand of insn I2 is from a retired instruction Ia.
- Because Ia is retired long ago, the physical destination register has been freed. The tag of the source register \$3 can not be renamed at the renaming stage for I2, still having the logical register tag \$3.
- Where does the operand \$3 of I2 come from?

Ia: add \$3,\$0,\$0  
 I1: sub p9,\$1,\$2  
 I2: add p10,p9,\$3  
 I3: or p11,\$4,\$5  
 I4: and p12,p10,p11



# Register renaming again

- A processor remembers a set of renamed logical registers.
- If  $\$1$  and  $\$2$  are not renamed for in-flight instructions, it uses  $p1$  and  $p2$  instead of  $\$1$  and  $\$2$ .



## Cycle 1

I1: sub \$5,\$1,\$2  
I2: add \$9,\$5,\$4  
I3: or \$5,\$5,\$2  
I4: and \$2,\$9,\$1

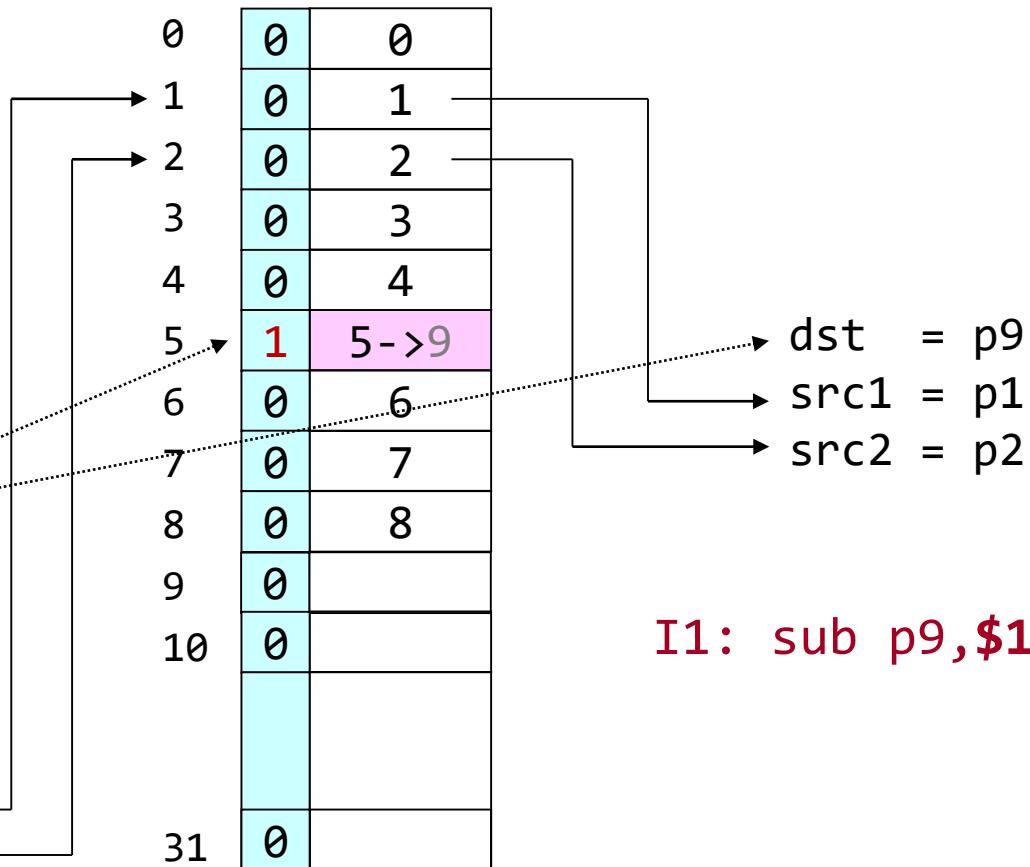
### Free tag buffer

	13	12	11	10	9
--	----	----	----	----	---

head

dst = \$5  
src1 = \$1  
src2 = \$2

Register map table



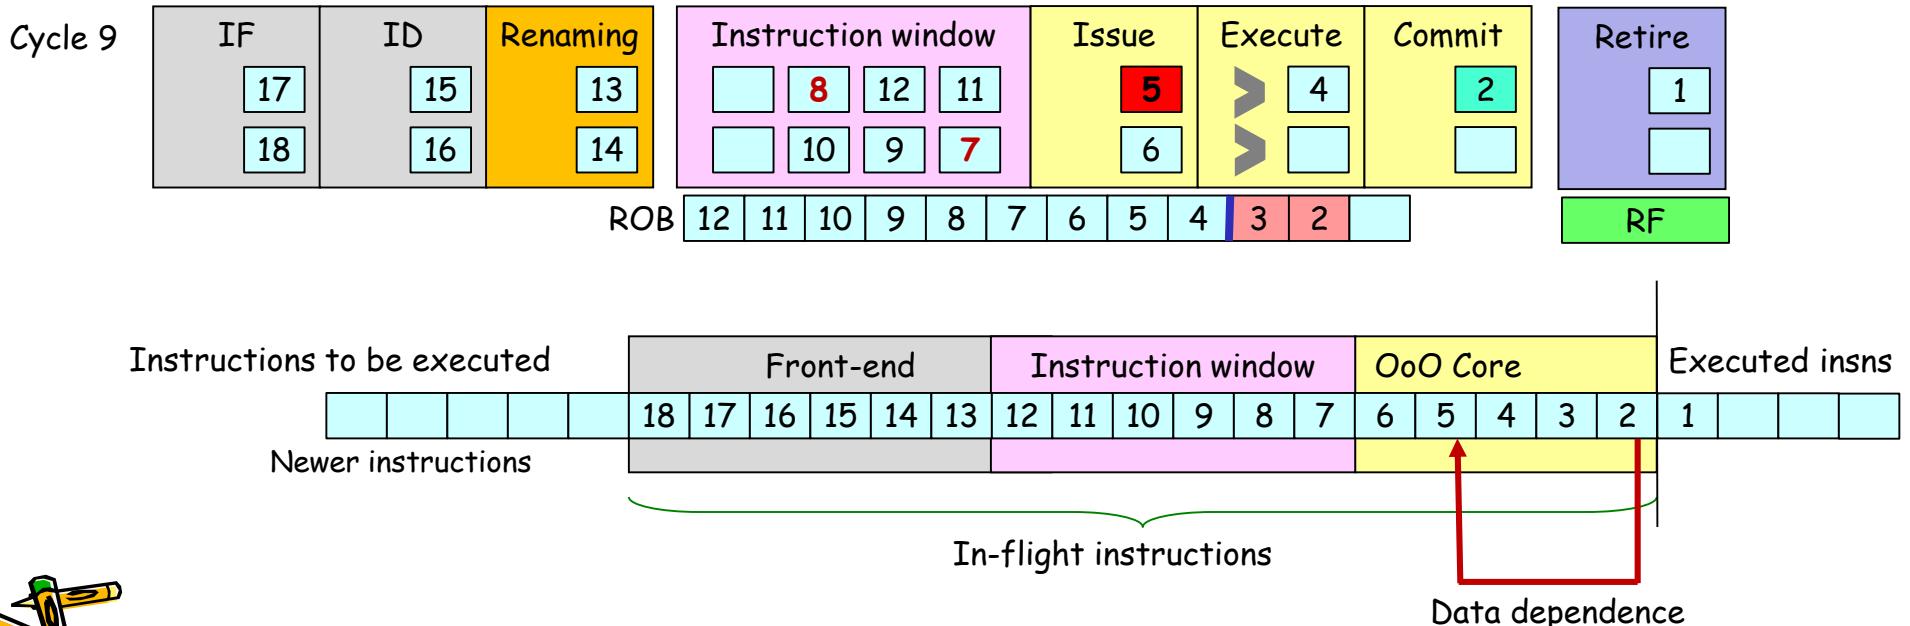
I1: sub p9,\$1,\$2



# Case 2: Register dataflow

- Assume that one source operand **p10** of insn I5 is from I2 which is not retired. The operand is generated a few clock cycles (tens of cycles sometimes) earlier.
- Because I2 is not retired, RF does not have the operand. Because I2 is committed, the operand is stored in ROB.
- Where does the operand of I5 come from?

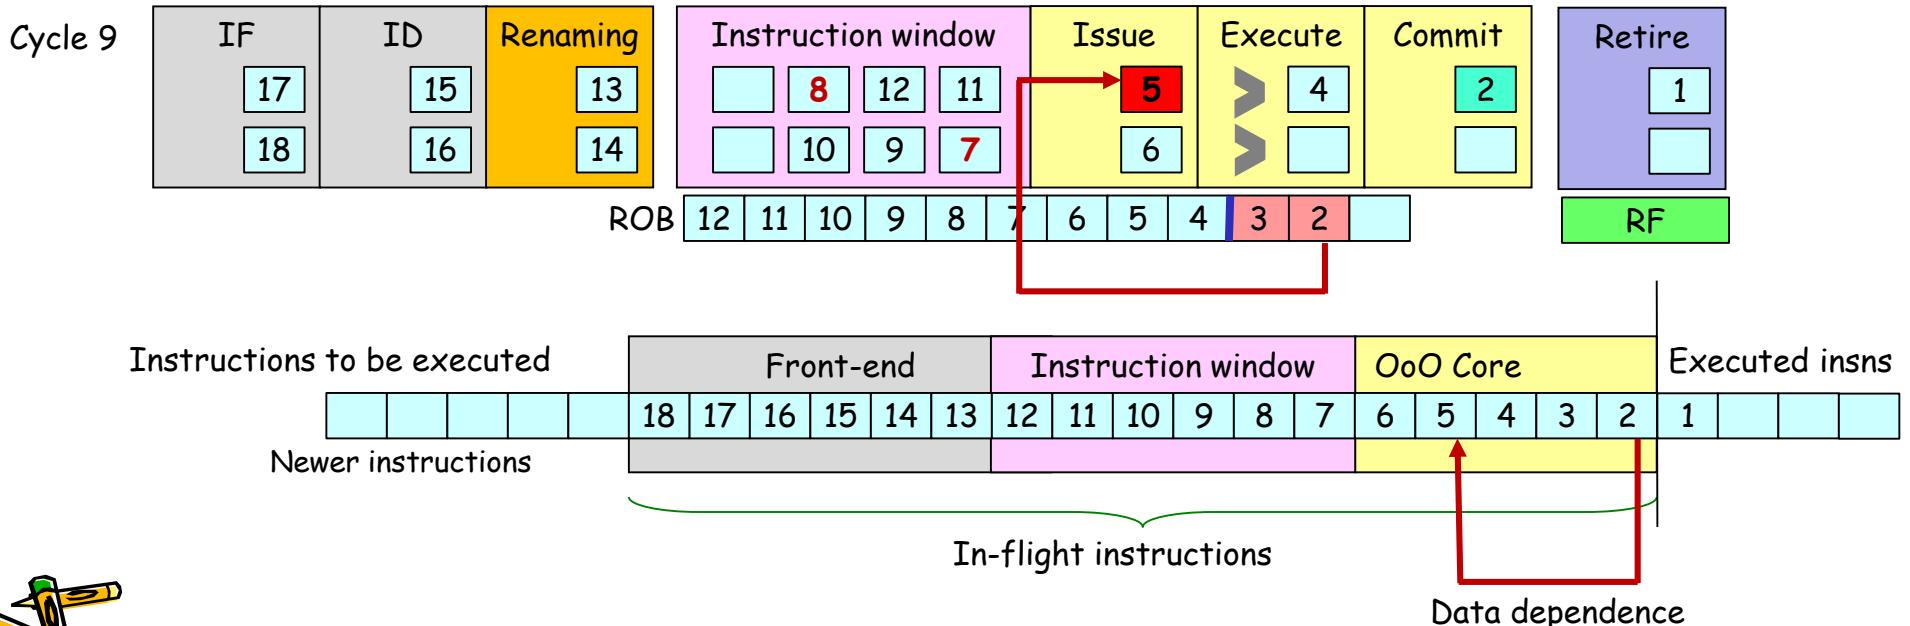
Ia: add \$3,\$0,\$0  
 I1: sub p9,\$1,\$2  
 I2: add p10,p9,\$3  
 I3: or p11,\$4,\$5  
 I4: and p12,p10,p11  
 I5: nor p13,p10,p12



# Case 2: Register dataflow from ROB

- Assume that one source operand **p10** of insn I5 is from I2 which is not retired. The operand is generated a few clock cycles (tens of cycles sometimes) earlier.
- Because I2 is not retired, RF does not have the operand. Because I2 is committed, the operand is stored in ROB.
- Where does the operand of I5 come from?

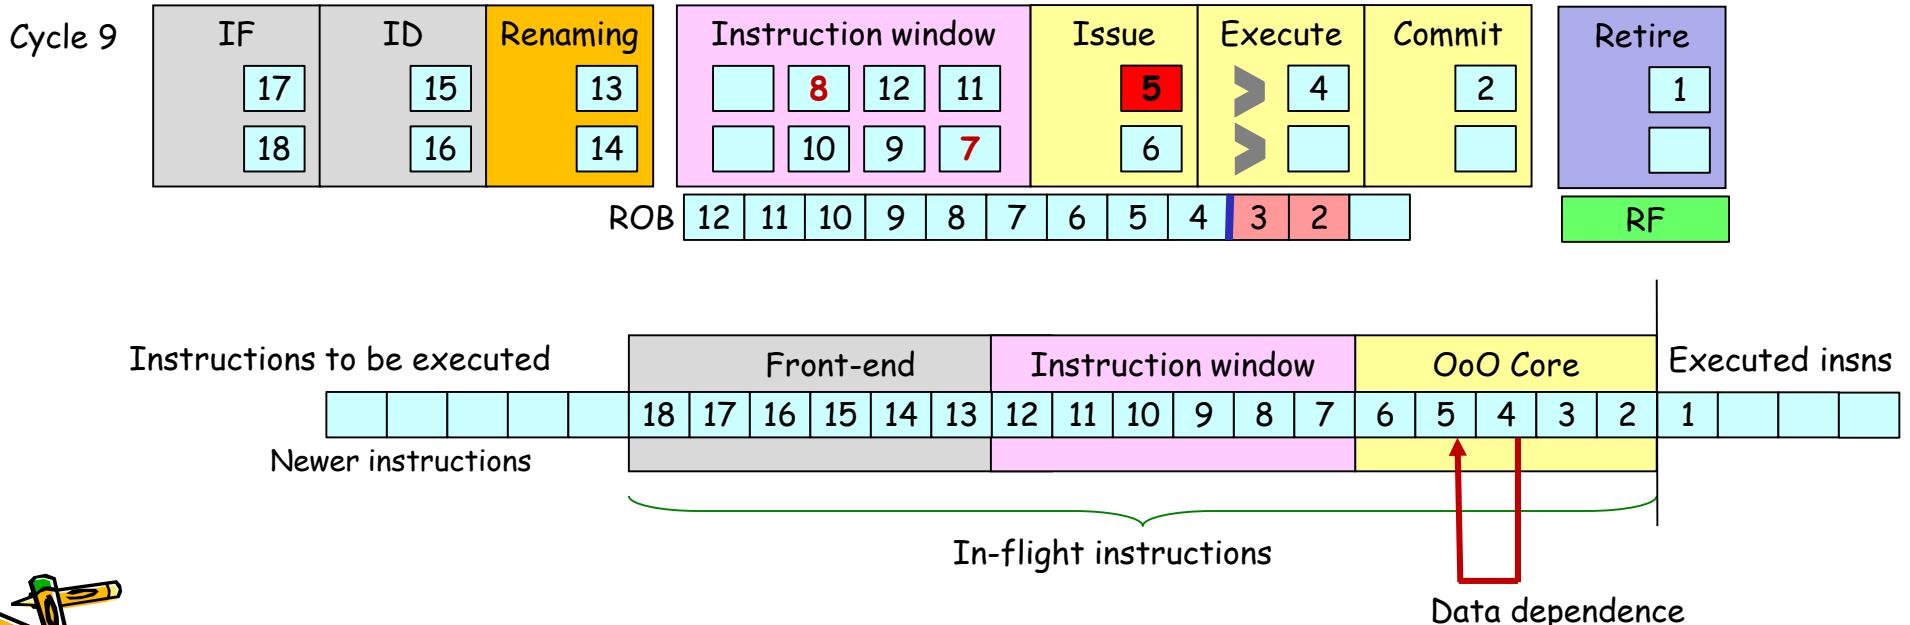
Ia: add \$3,\$0,\$0  
 I1: sub p9,\$1,\$2  
 I2: add p10,p9,\$3  
 I3: or p11,\$4,\$5  
 I4: and p12,p10,p11  
 I5: nor p13,p10,p12



# Case 3: Register dataflow

- Assume that the other source operand **p12** of insn I5 is from I4 which is not committed. The operand is generated in the previous clock cycle.
- Because I2 is not retired, RF does not have the operand.
- Because I2 is not committed, ROB does not have the operand.
- Where does the operand of I5 come from?

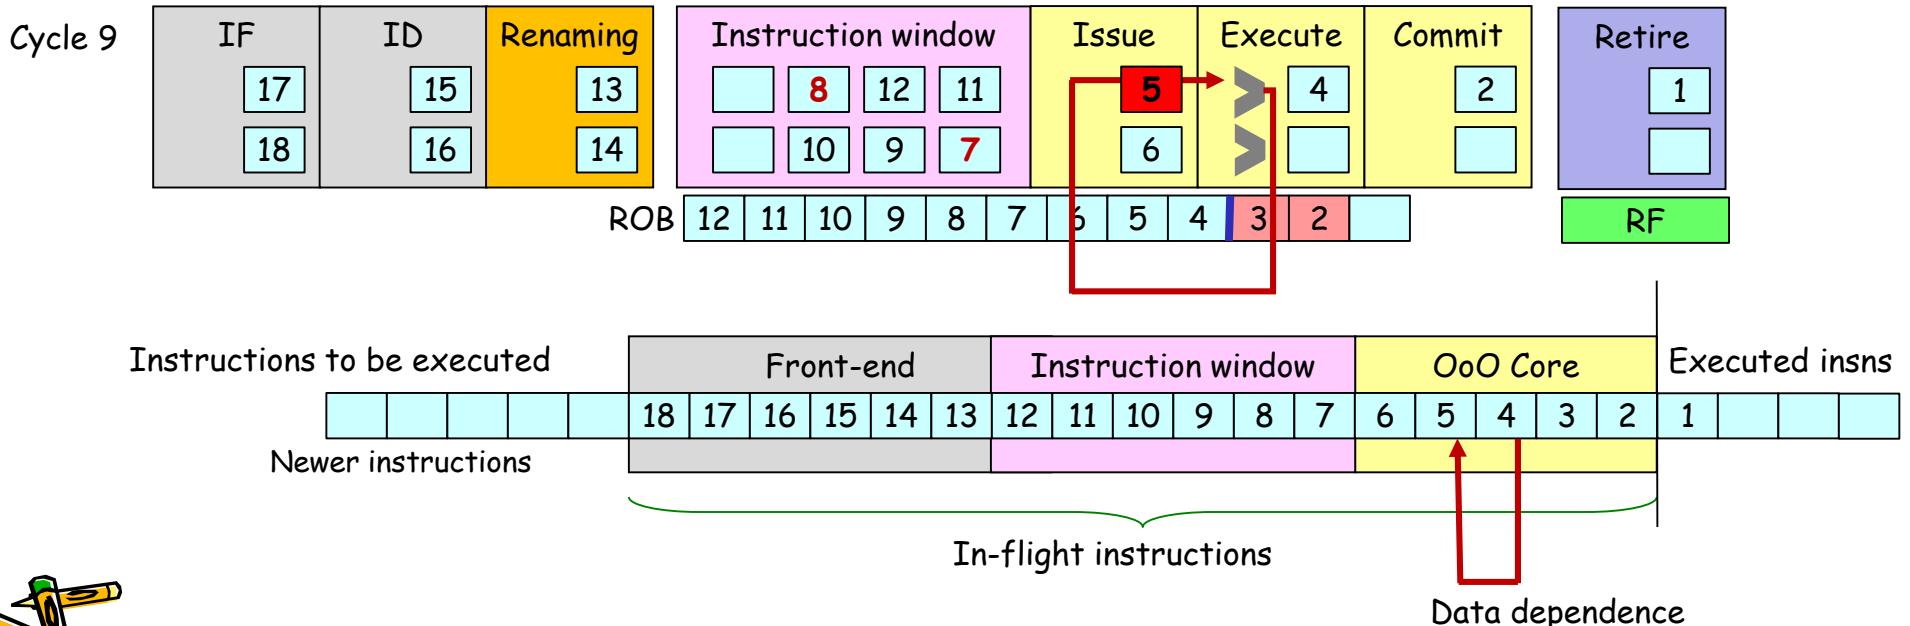
Ia: add \$3,\$0,\$0  
 I1: sub p9,\$1,\$2  
 I2: add p10,p9,\$3  
 I3: or p11,\$4,\$5  
 I4: and p12,p10,p11  
 I5: nor p13,p10,p12



# Case 3: Register dataflow from ALUs

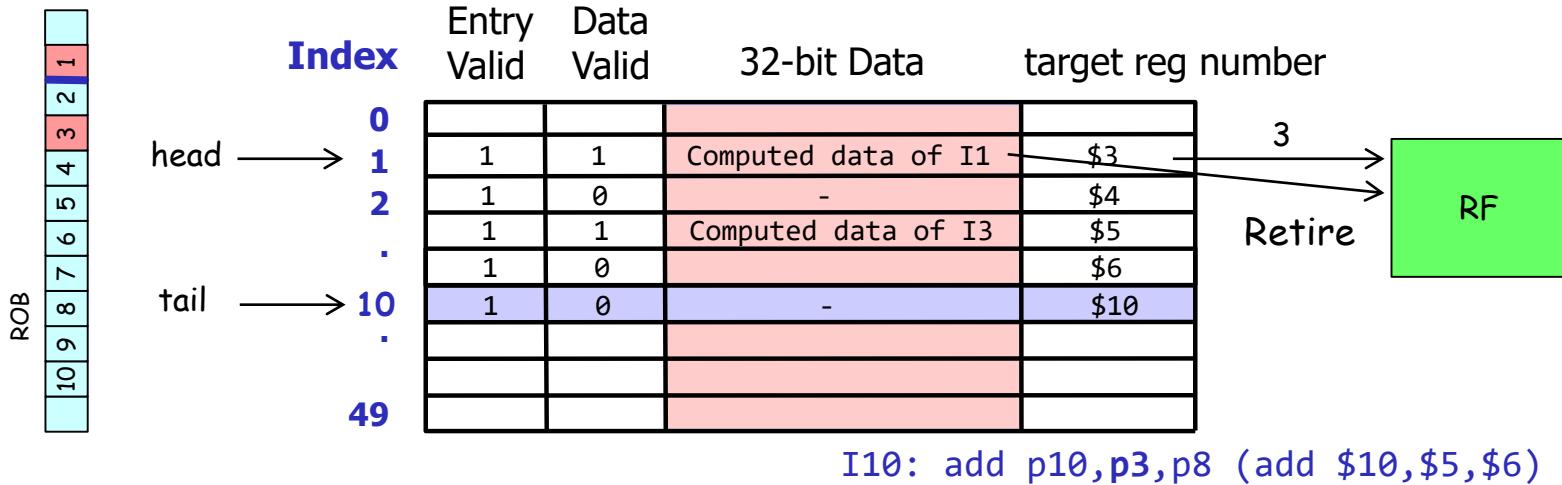
- Assume that the other source operand **p12** of insn I5 is from I4 which is not committed. The operand is generated in the previous clock cycle.
- Because I2 is not retired, RF does not have the operand.
- Because I2 is not committed, ROB does not have the operand.
- Where does the operand of I5 come from?

Ia: add \$3,\$0,\$0  
 I1: sub p9,\$1,\$2  
 I2: add p10,p9,\$3  
 I3: or p11,\$4,\$5  
 I4: and p12,p10,p11  
 I5: nor p13,p10,p12

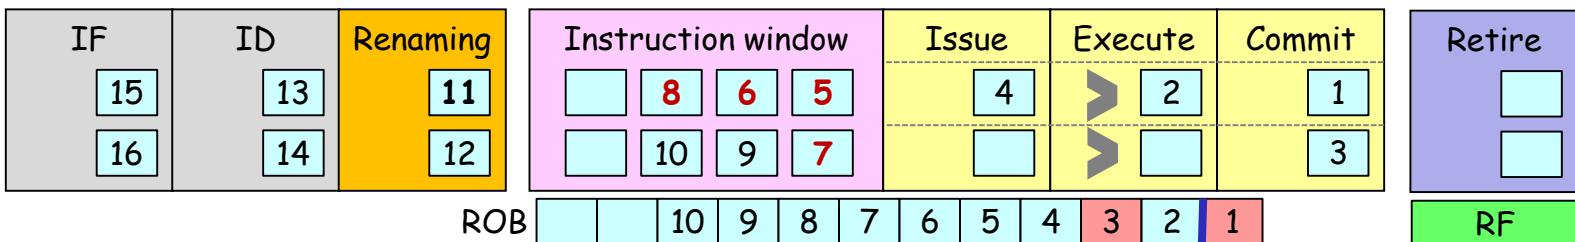


# Reorder buffer (ROB)

- Each ROB entry has following fields
  - entry valid bit, data valid bit, **data**, target register number, etc.
- ROB provides **the large physical registers** for renaming
  - in fact, physical register number is ROB entry number
- **The value of a physical register is from a matching ROB entry**



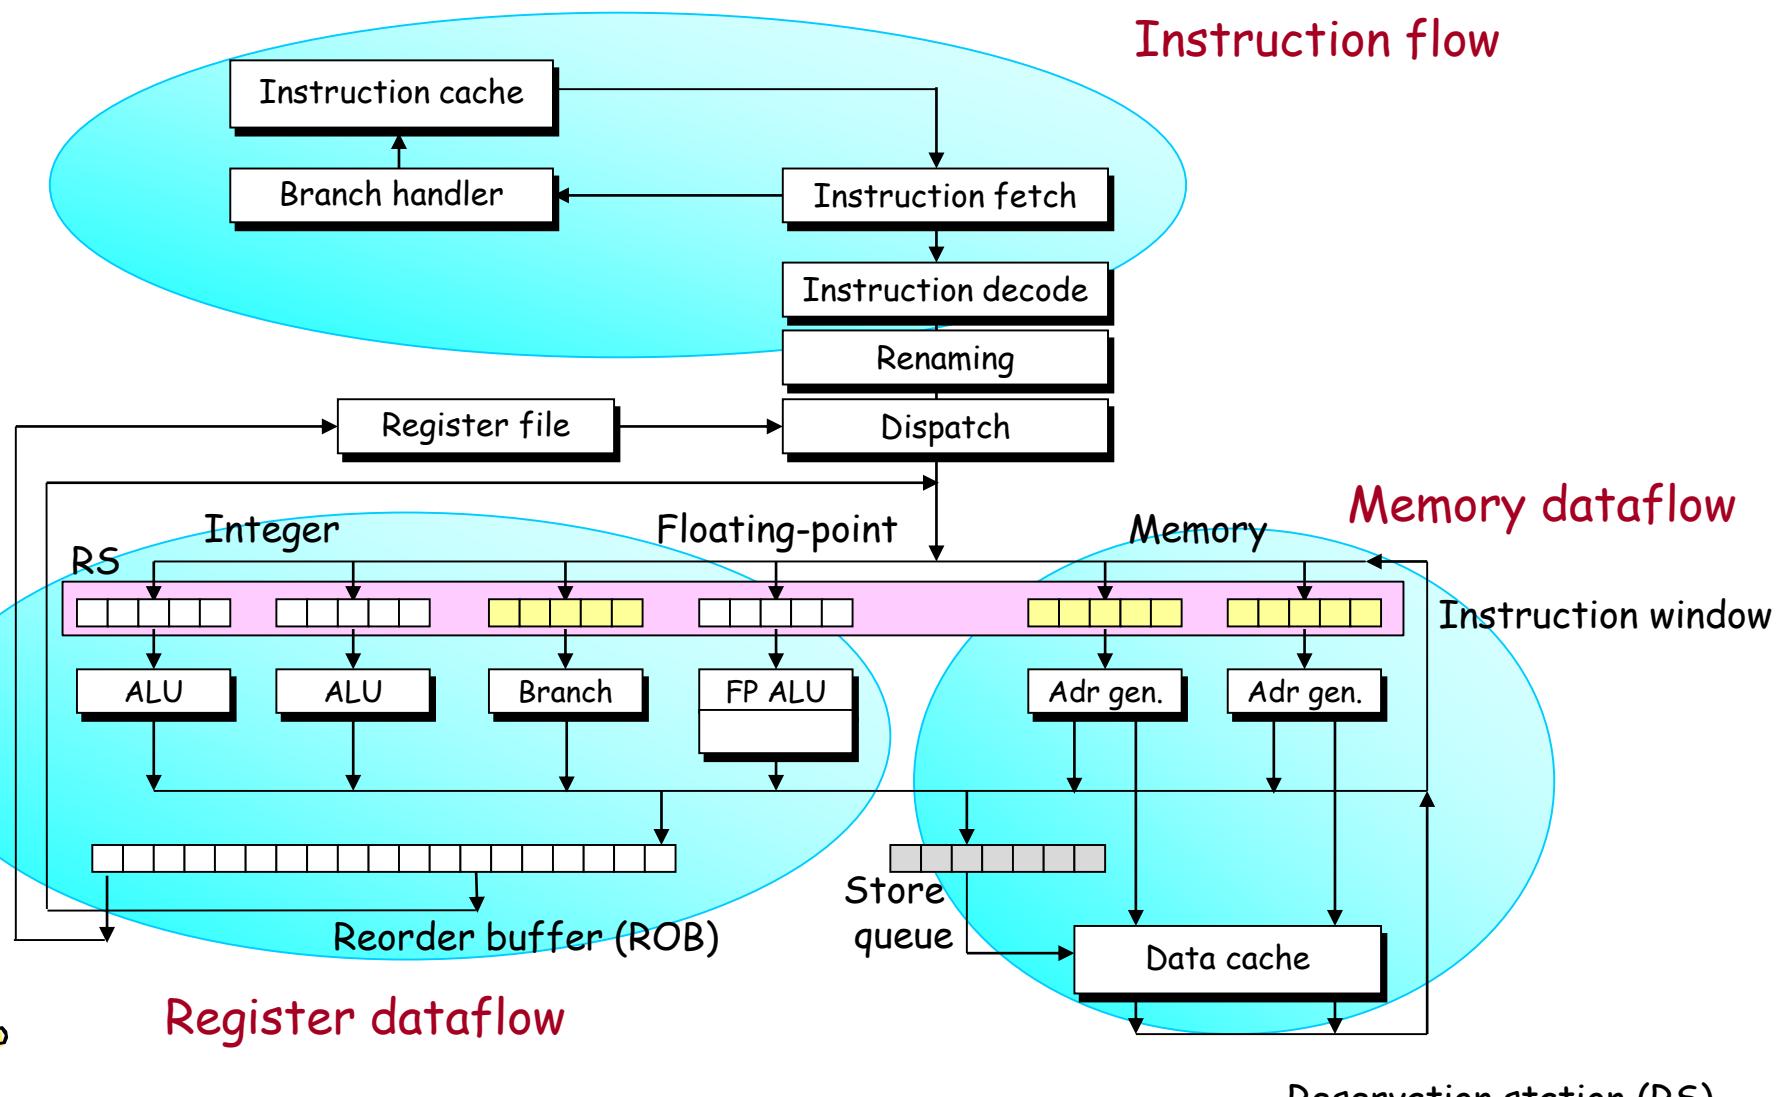
Cycle 8



this slide is to be used as a whiteboard

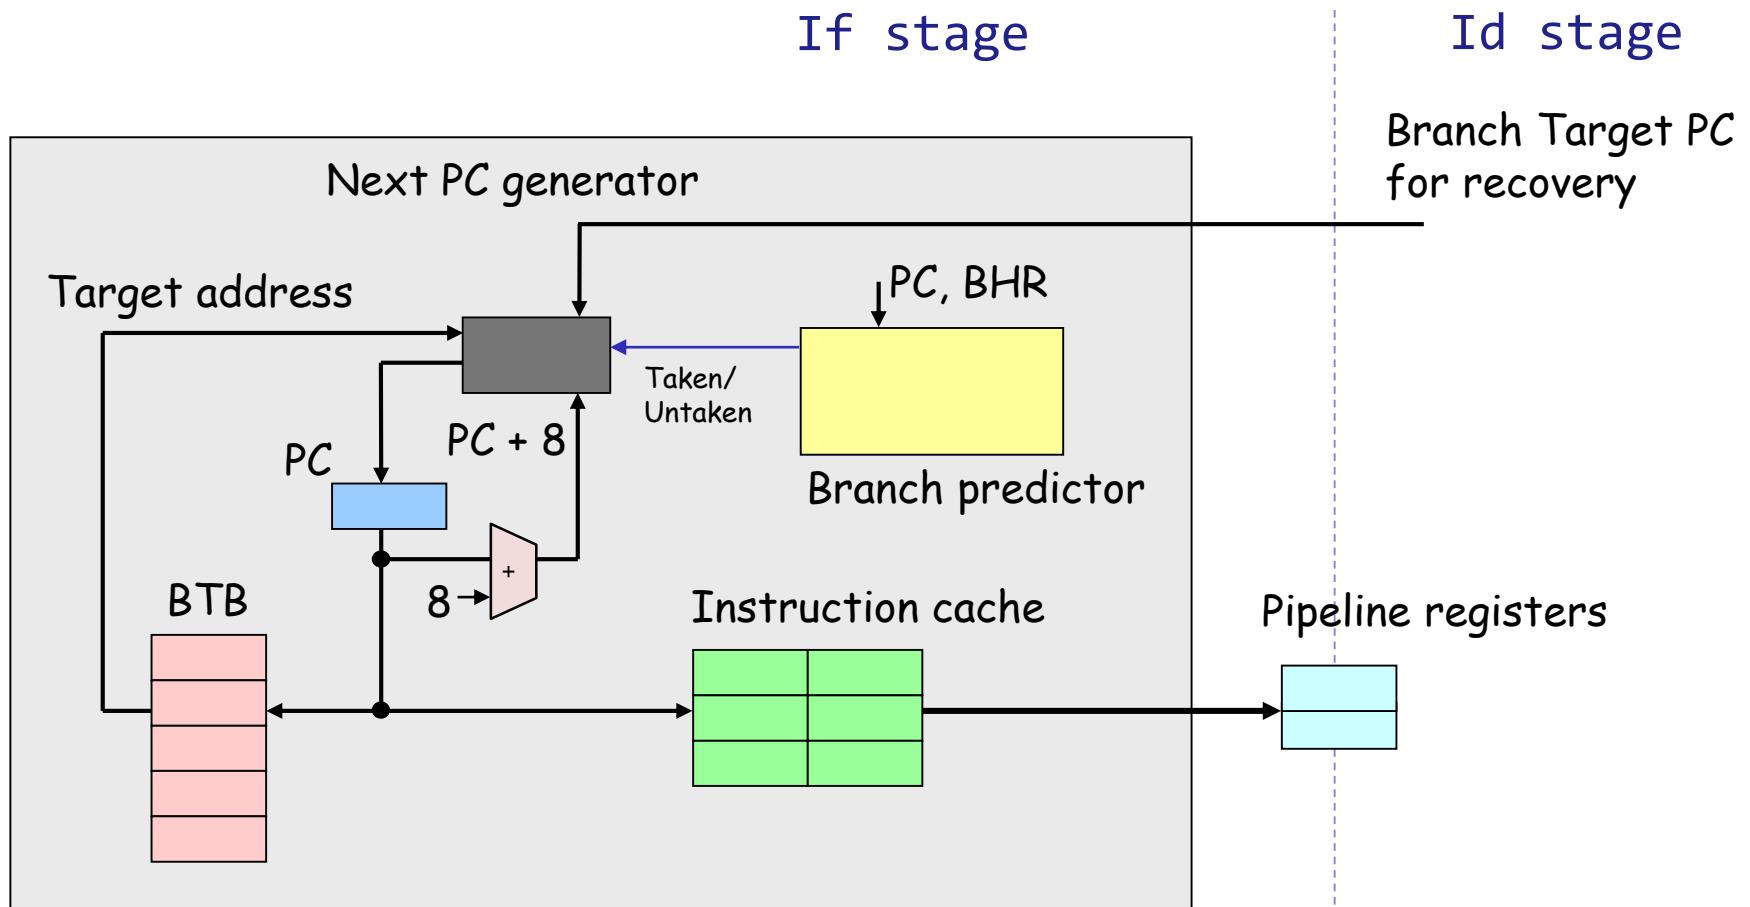


# Datapath of OoO execution processor



# Instruction fetch unit in IF stage

- For high-bandwidth instruction delivery, prediction, and speculation



# Renaming two instructions per cycle for superscalar

- Renaming instruction I0 and I1

Cycle 1

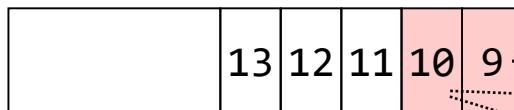
I0: sub \$5,\$1,\$2

I1: add \$9,\$5,\$4

I2: or \$5,\$5,\$2

I3: and \$2,\$9,\$1

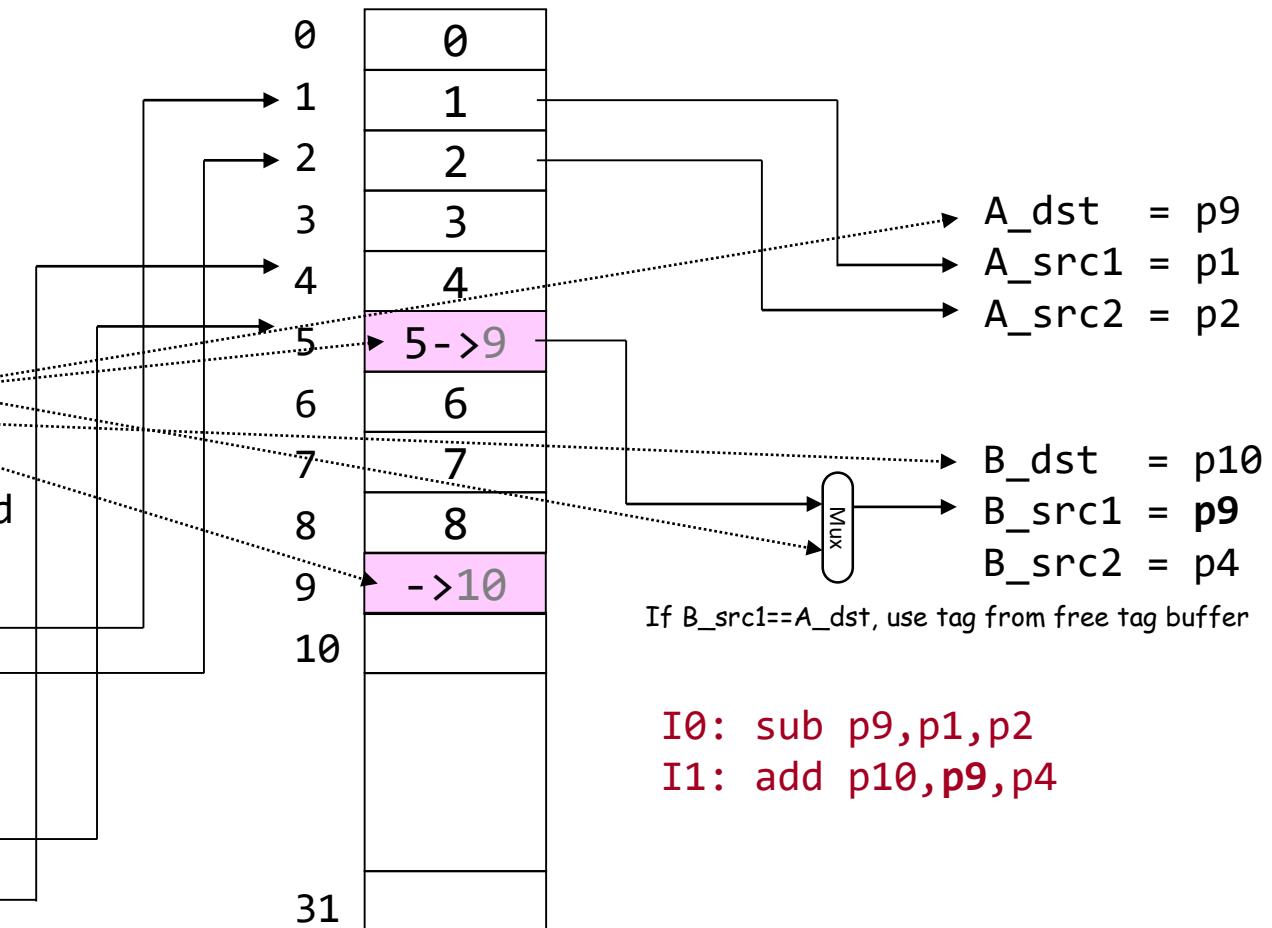
Free tag buffer



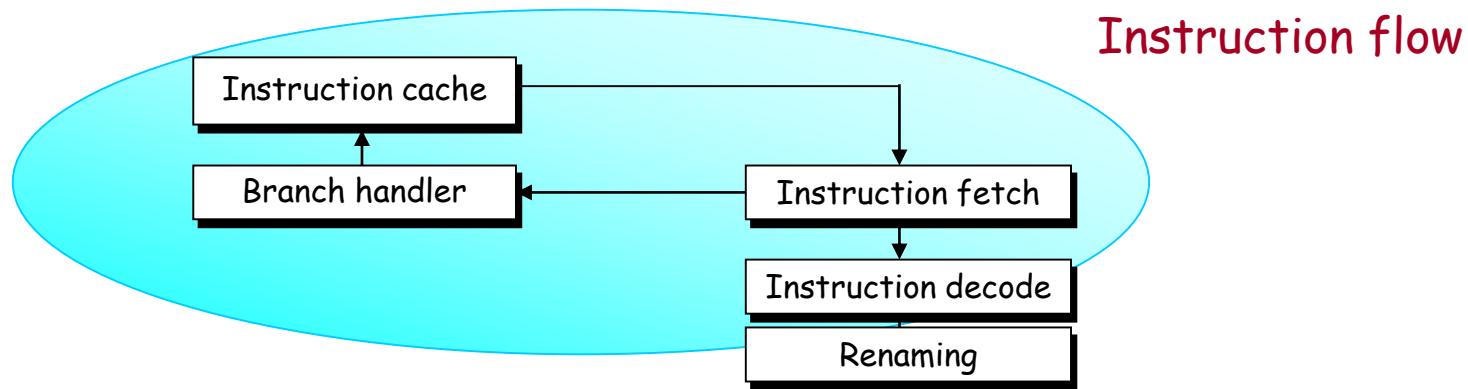
I0    **A\_dst** = \$5  
      A\_src1 = \$1  
      A\_src2 = \$2

I1    **B\_dst** = \$9  
      **B\_src1** = \$5  
      B\_src2 = \$4

Register map table

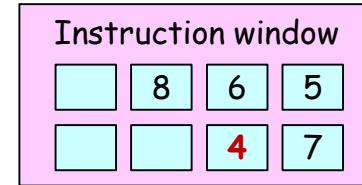
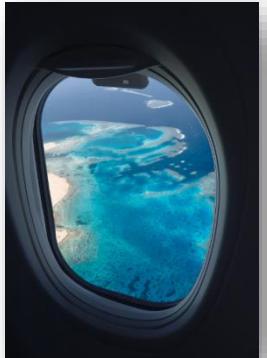


# Datapath of OoO execution processor (partially)

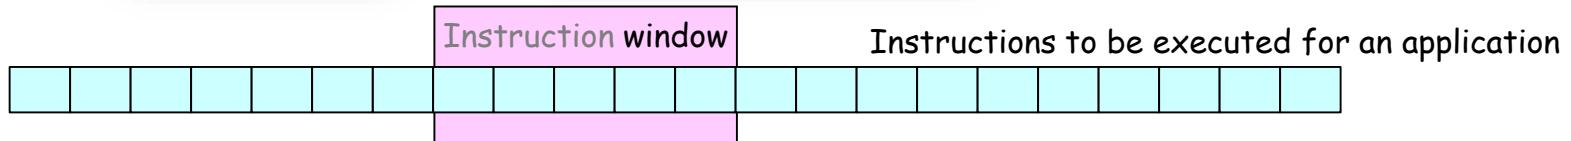


# Aside: What is a window?

- A window is a space in the wall of a building or in the side of a vehicle, which has glass in it so that light can come in and you can see out. (Collins)



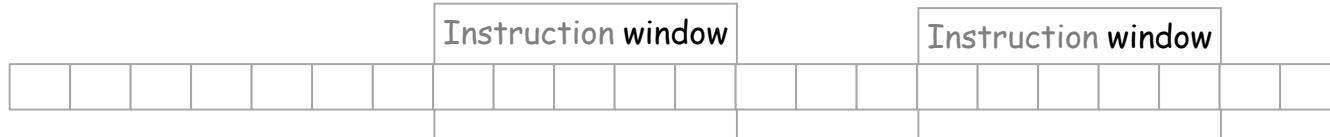
(a)



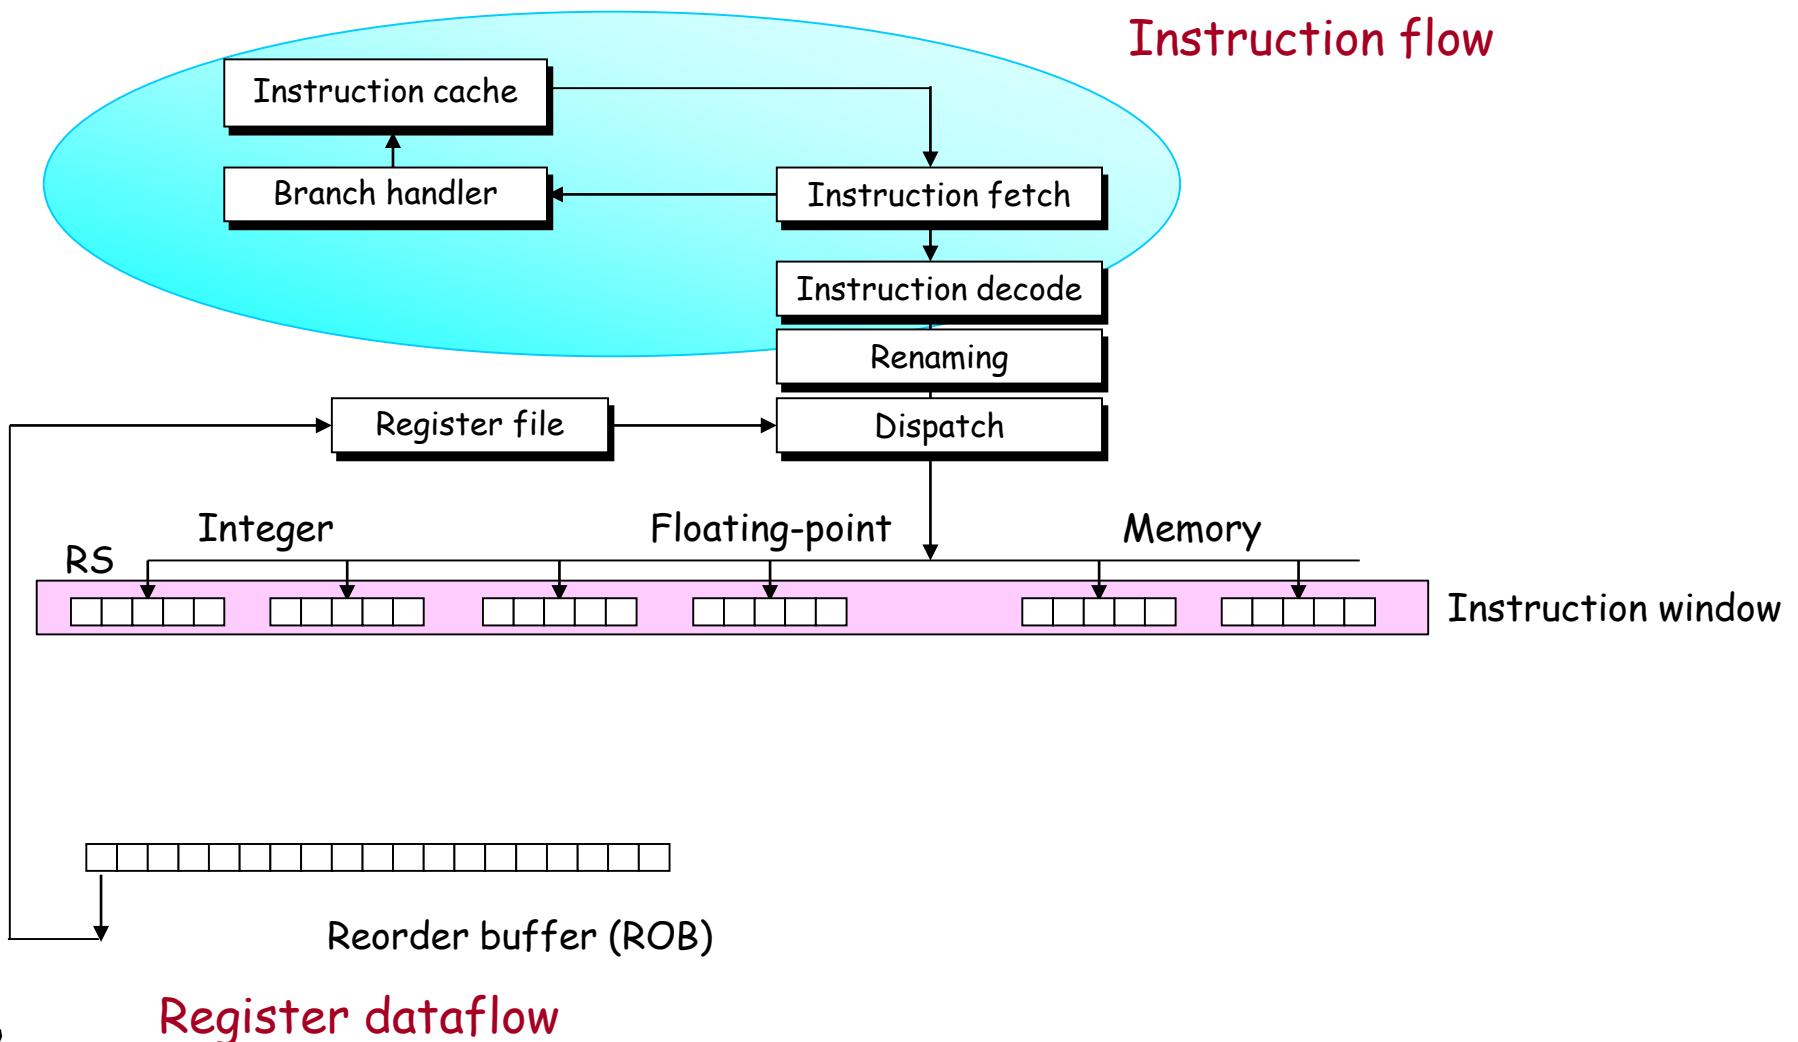
(b)



(c)



# Datapath of OoO execution processor (partially)

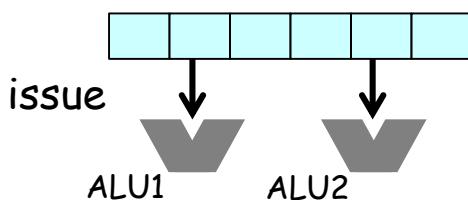


Register dataflow

# Reservation station (RS)

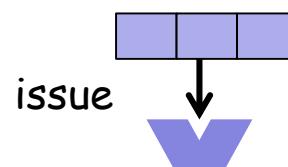
- To simplify the **wakeup** and **select** logic at issue stage, each functional unit (ALU) has own instruction window, an entry for an instruction is called **reservation station (RS)**.
- Each reservation station has
  - valid bit, src1 tag, src1 data, src1 ready, src2 tag, src2 data, src2 ready, destination physical register number (dst), operation, ...
  - The computed data with its *dst* as tag is broadcasted to all RSs.

instruction window for ALU1 and ALU2

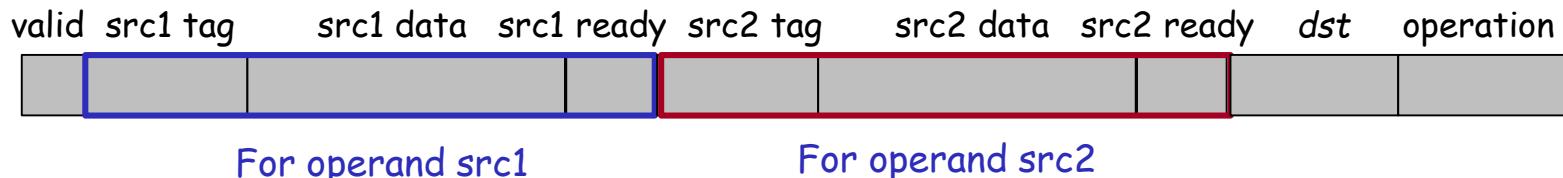


(a) Central instruction window

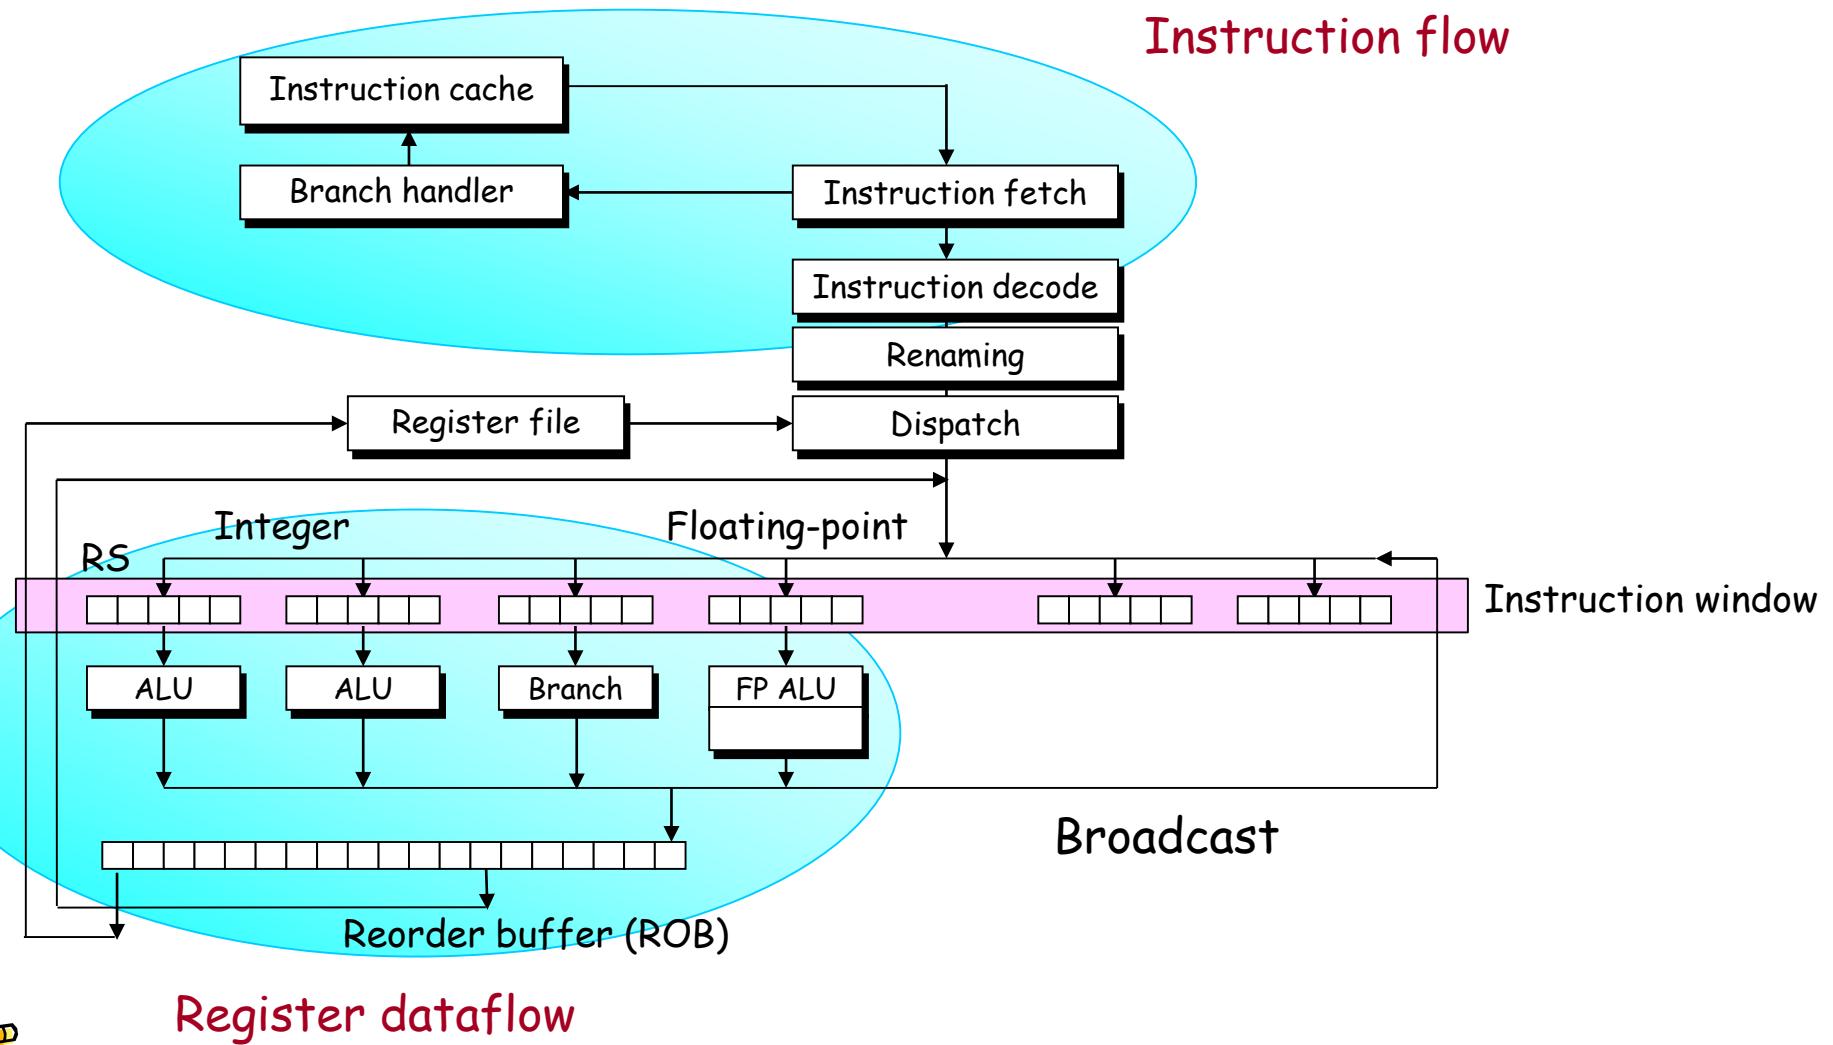
IW for ALU1



(b) instruction window using RS

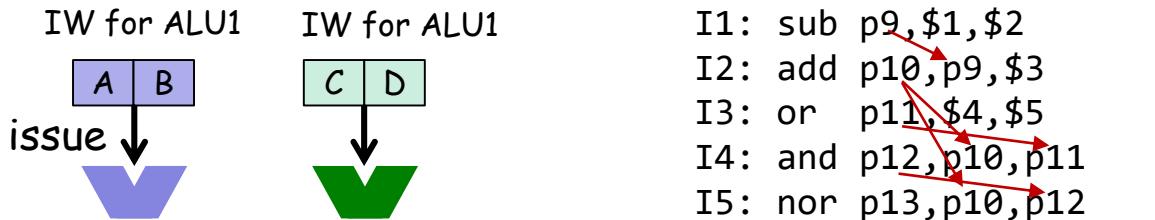


# Datapath of OoO execution processor (partially)

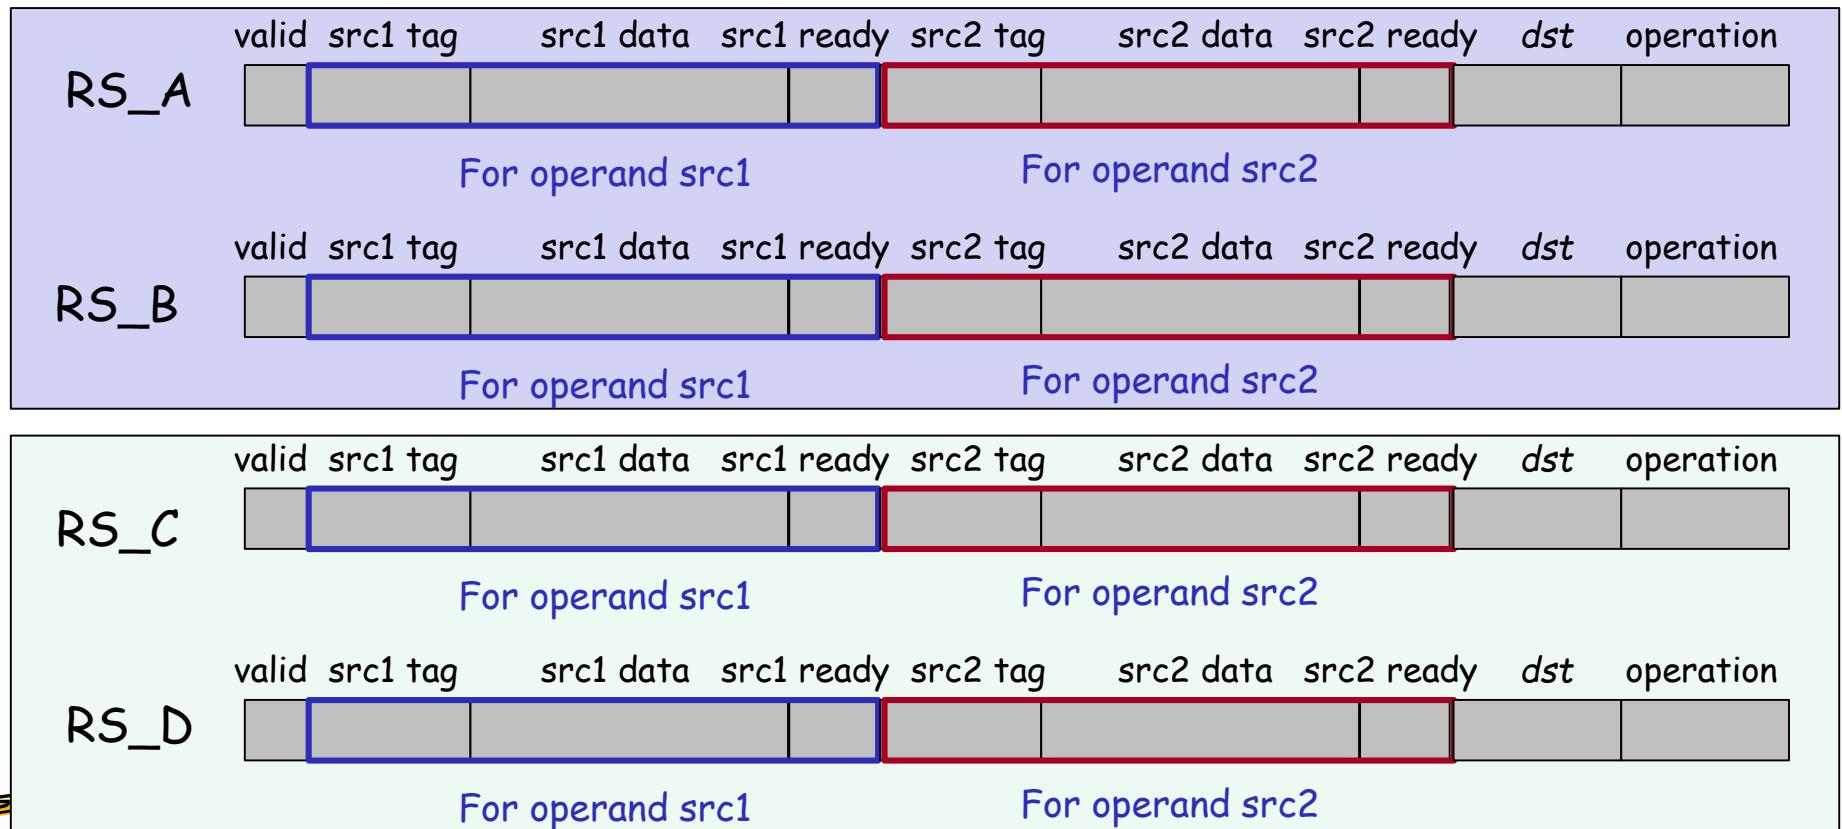


# Example behavior of reservation stations

Cycle 0

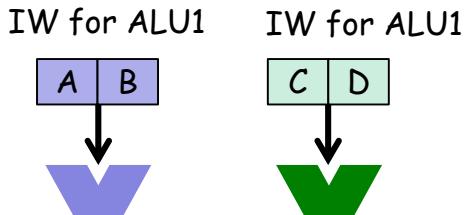


dispatch at most two instructions, one to A or B and the other to C or D



# Example behavior of reservation stations

Cycle 1



I1: sub p9,\$1,\$2  
I2: add p10,p9,\$3  
I3: or p11,\$4,\$5  
I4: and p12,p10,p11  
I5: nor p13,p10,p12

dispatch at most two instructions, one to A or B and the other to C or D

	valid	src1 tag	src1 data	src1 ready	src2 tag	src2 data	src2 ready	dst	operation
RS_A	1	\$1	value of \$1	1	\$2	value of \$2	1	p9	I1: sub
	For operand src1				For operand src2				
	valid	src1 tag	src1 data	src1 ready	src2 tag	src2 data	src2 ready	dst	operation
RS_B									
	For operand src1				For operand src2				

	valid	src1 tag	src1 data	src1 ready	src2 tag	src2 data	src2 ready	dst	operation
RS_C	1	p9		0	\$3	value of \$3	1	p10	I2: add
	For operand src1				For operand src2				
	valid	src1 tag	src1 data	src1 ready	src2 tag	src2 data	src2 ready	dst	operation
RS_D									
	For operand src1				For operand src2				

# Example behavior of reservation stations

Cycle 2

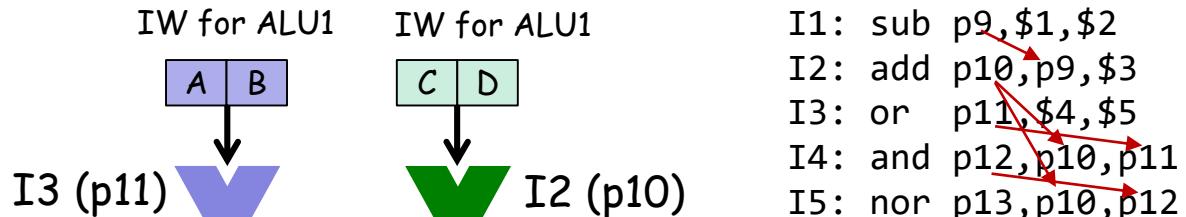


dispatch at most two instructions, one to A or B and the other to C or D

	valid	src1 tag	src1 data	src1 ready	src2 tag	src2 data	src2 ready	dst	operation
RS_A	0								
	For operand src1				For operand src2				
RS_B	1	\$4	value of \$4	1	\$5	value of \$5	1	p11	I3: or
	For operand src1				For operand src2				
RS_C	1	p9	value of p9	1	\$3	value of \$3	1	p10	I2: add
	For operand src1				For operand src2				
RS_D	1	p10		0	p11		0	p12	I4: and
	For operand src1				For operand src2				

# Example behavior of reservation stations

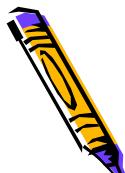
Cycle 3



dispatch at most two instructions, one to A or B and the other to C or D

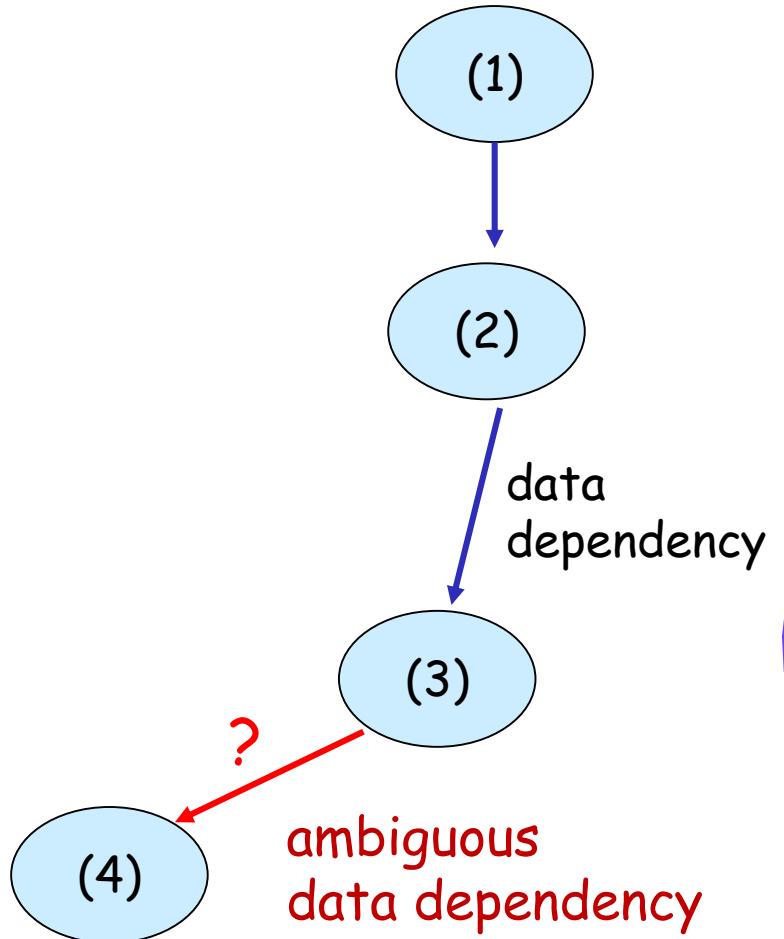
	valid	src1 tag	src1 data	src1 ready	src2 tag	src2 data	src2 ready	dst	operation
RS_A	1	p10	value of p10	1	p12			0	p13 I5: nor
	For operand src1				For operand src2				
RS_B	0								
	For operand src1				For operand src2				
RS_C	1								I6:
	For operand src1				For operand src2				
RS_D	1	p10	value of p10	1	p11	value of p11	1	p12	I4: and
	For operand src1				For operand src2				

this slide is to be used as a whiteboard



# Instruction Level Parallelism (ILP)

lw	$\$t0, 32(\$s3)$	(1)
add	$\$t0, \$s2, \$t0$	(2)
sw	$\$t0, 48(\$s3)$	(3)
lw	$\$t1, 32(\$s4)$	(4)

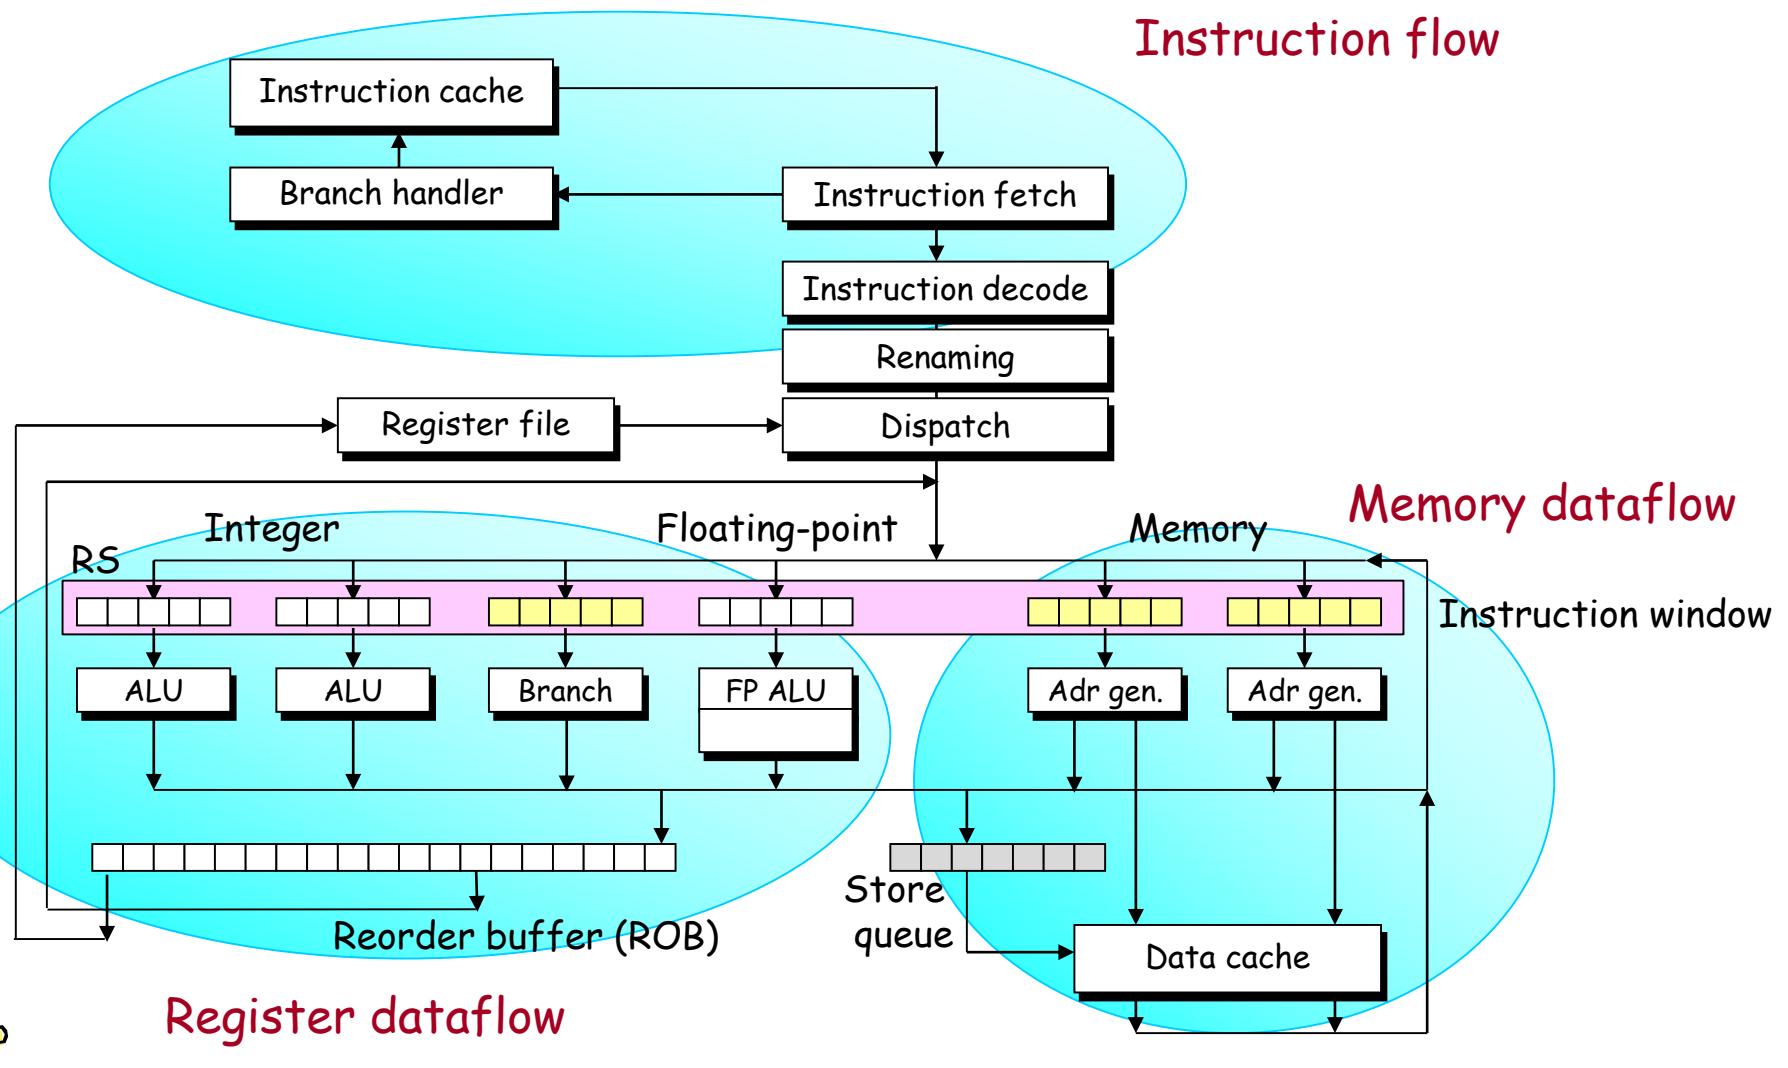


# Memory dataflow and branches

- The update of a data cache cannot be recovered easily. So, **cache update is done at the retire stage** in-order manner by using **store queue**.  
Because of the **ambiguous memory dependency**, load and store **instructions** can be executed in-order manner.
  - About 30% (or less) of executed instructions are load and stores.
  - Even if they are executed in-order, IPC of 3 can be achieved.
- **Branch instructions** can be executed in-order manner.
  - About 20% (or less) of executed instructions are jump and branch instructions.
  - Out-of-order branch execution and aggressive miss recovery may cause false recovery (recovery by a branch on the false control path).

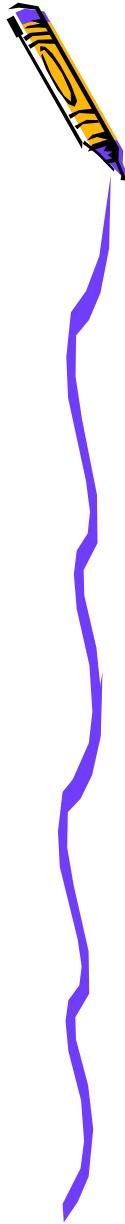


# Datapath of OoO execution processor



# Pollack's Rule

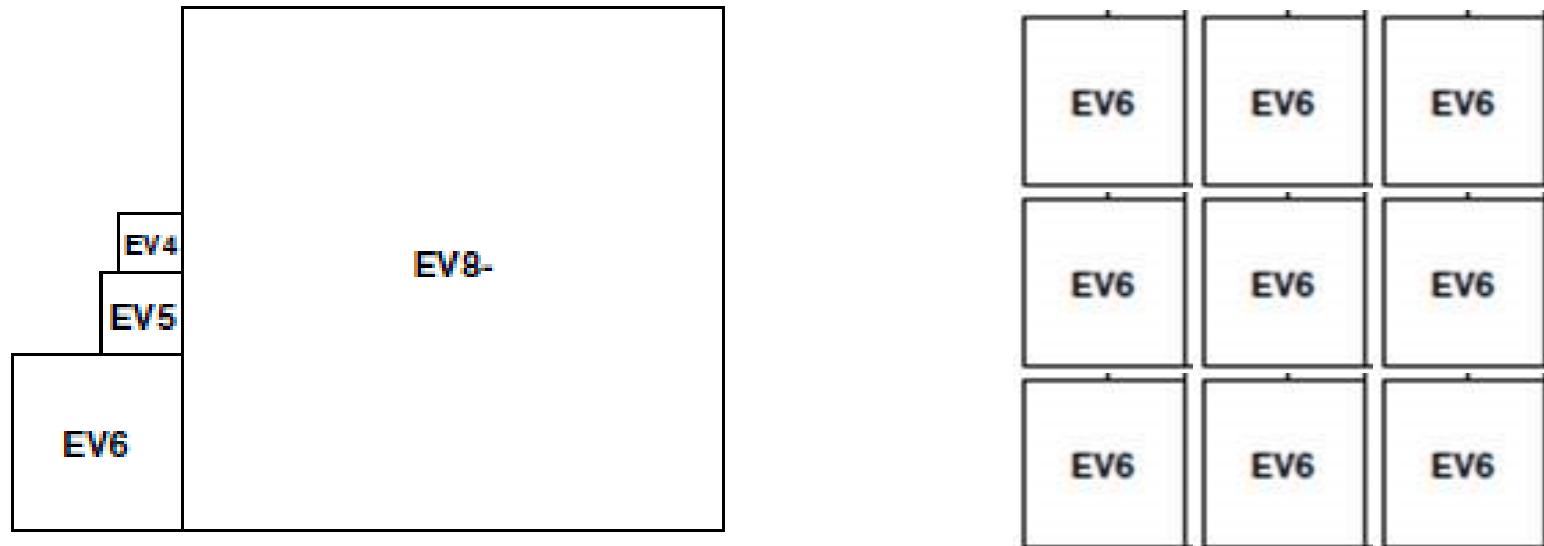
- Pollack's Rule states that microprocessor "performance increase due to microarchitecture advances is roughly proportional to the square root of the increase in complexity". Complexity in this context means processor logic, i.e. its area.



WIKIPEDIA



# From multi-core era to many-core era



**Figure 1. Relative sizes of the cores used in the study**

Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction, MICRO-36

# From multi-core era to many-core era

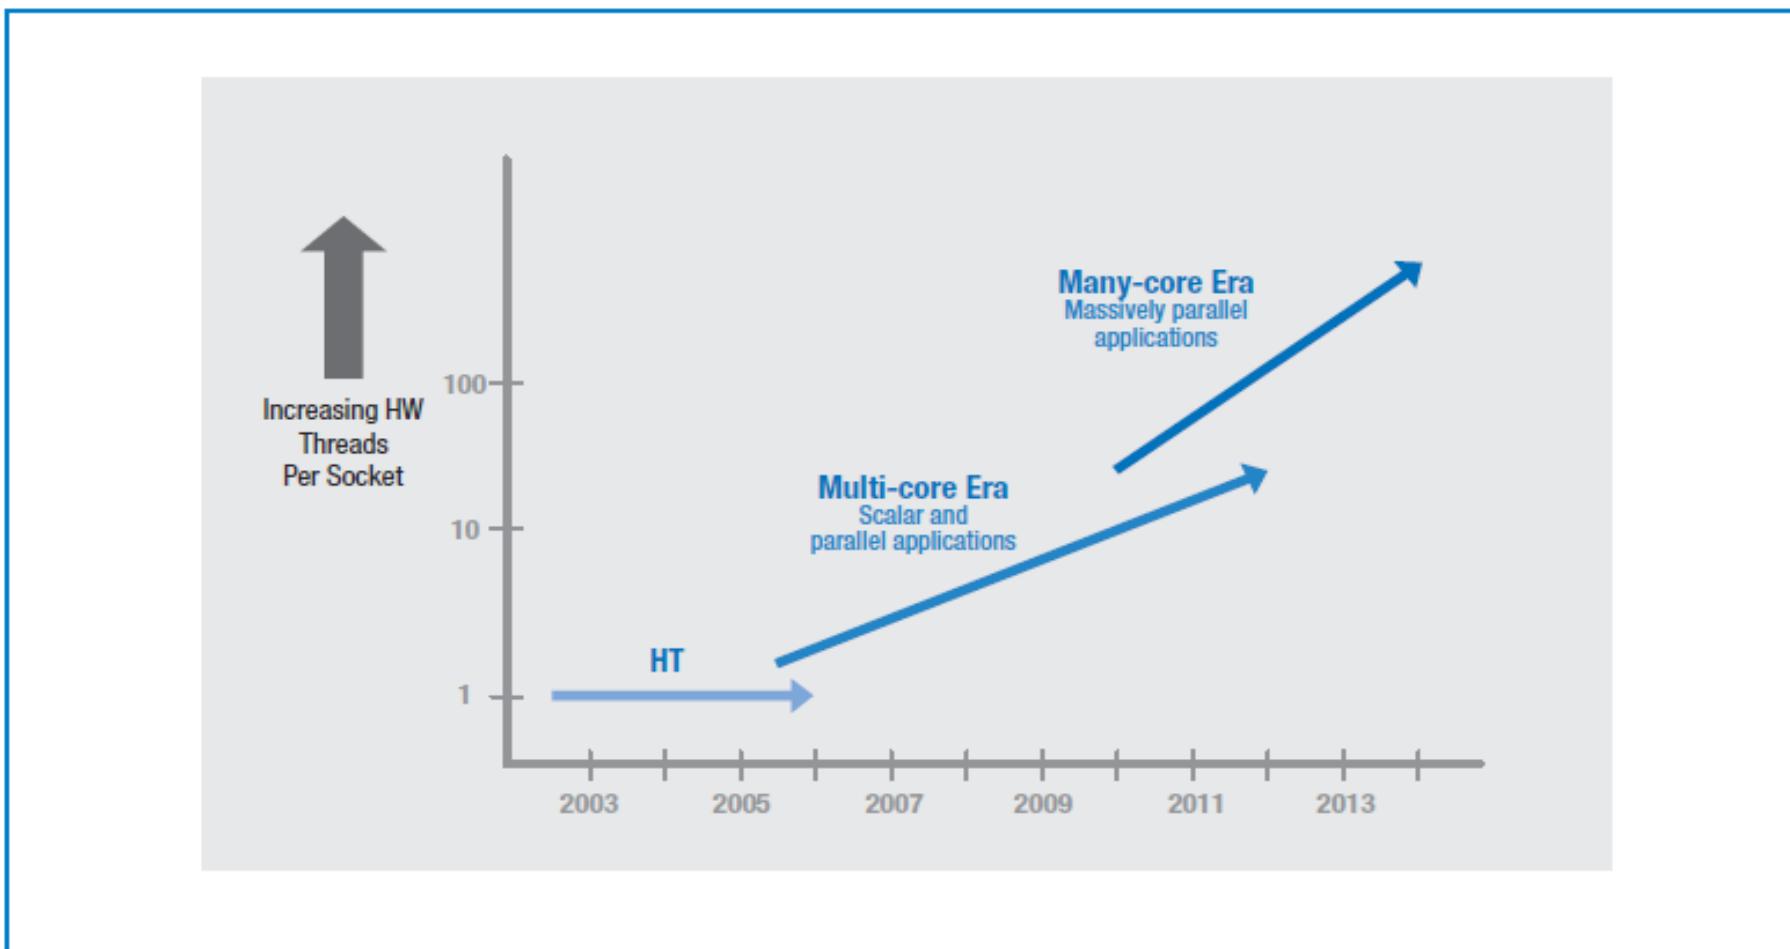


Figure 1: Current and expected eras of Intel® processor architectures

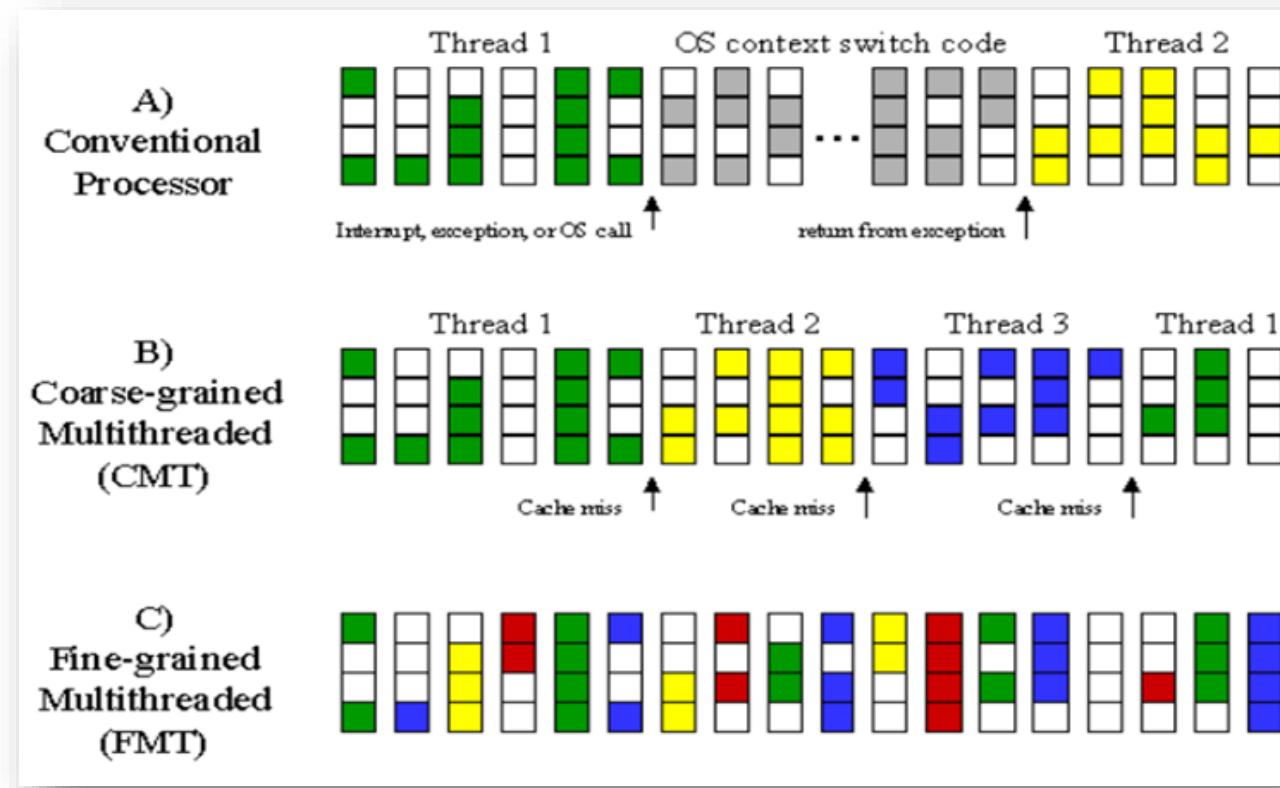
Platform 2015: Intel® Processor and Platform Evolution for the Next Decade, 2005

this slide is to be used as a whiteboard



# Multithreading (1/2)

- During a branch miss recovery and access to the main memory by a cache miss, ALUs have no jobs to do and have to be idle.
- Executing **multiple independent threads (programs)** will mitigate the overhead.
- They are called coarse- and fine-grained multithreaded processors having multiple architecture states.



# Multithreading (2/2)

- Simultaneous Multithreading (SMT) can improve hardware resource usage.

